




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 13/00	A1	(11) International Publication Number: WO 98/02821 (43) International Publication Date: 22 January 1998 (22.01.98)
--	-----------	--

(21) International Application Number: PCT/US97/12314
(22) International Filing Date: 15 July 1997 (15.07.97)
(30) Priority Data: 08/683,570 15 July 1996 (15.07.96) US
(71) Applicant: 3COM CORPORATION [US/US]; 5400 Bayfront Plaza, P.O. Box 58145, Santa Clara, CA 95052-8145 (US).
(72) Inventors: HART, John, H.; 15121 Pepper Lane, Saratoga, CA 95070 (US). SHERER, W., Paul; 1054 Gardenia Way, Sunnyvale, CA 94086 (US).
(74) Agent: RICHARDSON, Kent, R.; Wilson Sonsini Goodrich & Rosati, 650 Page Mill Road, Palo Alto, CA 94304-1050 (US).

(81) Designated States: AU, CA, GB, JP, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).
Published
With international search report.

(54) Title: VIRTUAL NETWORK ARCHITECTURE FOR CONNECTIONLESS LAN BACKBONE



Best Available Copy

(57) Abstract

Network traffic management is achieved based on automatically setting up a plurality (127, 128) of virtual networks (VNETs) within a single large virtual LAN. Multicast/broadcast traffic is confined to the VNET (127, 128) of the source, without imposing constraints on layer two addressing within the virtual LAN. VNETs are domains of users of a virtual LAN which include members of logical networks defined at layer three or higher. One method includes transferring a multi-destination packet originating from a particular node in the virtual LAN by tunnelling across a connectionless backbone network (120) to a virtual net server (125). The virtual net server (125) translates the multi-destination packet to a plurality of tunneled messages identifying nodes authorized to receive multi-destination packets from the members of the particular VNET which originated the packet. The tunneled messages are then forwarded from the virtual net server to the authorized nodes.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

VIRTUAL NETWORK ARCHITECTURE FOR
CONNECTIONLESS LAN BACKBONE

CONTINUING APPLICATION DATA

5 The present application is a continuation-in-part of copending U.S. Patent Application No. 08/502,835; filed July 14, 1995; invented by John H. Hart; entitled VIRTUAL NETWORK ARCHITECTURE.

BACKGROUND OF THE INVENTION

10 Field of the Invention

 The present invention relates to data communication networks which consist of a number of local area network (LAN) segments interconnected to form a virtual LAN environment; and more particularly to methods for managing data flow in such networks across a connectionless LAN backbone.

15

Background of the Invention

 Historically, networks have been designed around the wired LAN segment as the basic technique for establishing network user groups. Standard network layer protocols define logical networks with a single layer two (data link layer) LAN segment in mind, with layer two bridging and layer three (network layer) routing functions used for moving data between LAN segments and layer three logical networks. However, with the emerging ATM LAN emulation mode and other LAN switching systems, the layer two boundaries become less controlled, giving rise to the concept of a virtual LAN. See, United States Patent No. 20 4,823,338 to Chan *et al.*, and an IEEE standard referred to as 802.1D. Nodes in a single layer two virtual LAN are found on different physical LAN segments but have the appearance to layer two processes (data link layer processes using medium access control MAC addresses) of residing on a single layer two LAN segment. This allows a unicast packet to propagate across the virtual LAN to 25 any other station in the virtual LAN. Also, multi-destination packets generated 30

on a particular LAN segment propagate throughout a number of interconnected LAN segments to ensure that all possible members of the virtual LAN receive the packet.

5 Within virtual LAN domains, multicast/broadcast frames are used by higher layer "discovery" or "advertisement" procedures to locate other systems or services within the virtual LAN domain. Systems send "data" to other systems using unicast MAC address which are either known in advance or learned through multicast /broadcast discovery and advertisement procedures. Systems send "multi-media data" using either unicast or multicast frames with
10 special protocols to improve throughput or latency, as required.

 Large virtual LANs create large multicast/broadcast domains; and the burden on the backbone network of transmitting all these multi-destination packets begins to impact overall system performance. More importantly, the users of the virtual LAN become burdened by a large number of multi-
15 destination packets that must be inspected and processed, even when the packet is simply discarded. In fact, several layer three network protocols may co-exist in a single virtual LAN, resulting in much traffic which is irrelevant to many users in the virtual LAN, which must nonetheless process the traffic to discover that the network layer data unit carried in it relates to a protocol it does not use.

20 Commonly used network layer protocols include the internet protocol (IP) originally developed under DARPA, the interpacket exchange protocol (IPX) published by Novell, the Xerox network system (XNS) published by Xerox, the Banyan VINES protocol, the NetBIOS protocol published by IBM and Microsoft, AppleTalk published by Apple Computer, and the DECNet protocol
25 published by Digital Equipment Corporation. Many network layer protocols create protocol specific domains based on the logical network identifiers. For example, the IP protocol establishes "subnet" domains based on the network number portion, and extensions, of the IP address of the frame. The IPX protocol creates logical networks based on the internal network number
30 assigned to servers in the network. AppleTalk creates "zones". The NetBIOS

protocol does not support multiple domains within a single LAN or emulated LAN, and can thus be considered to define a single (or "null") logical network at layer three, by default. These protocol specific logical networks defined at layer three, or higher layers, are called virtual networks, or VNETs in the present application. By the nature of virtual LANs according the prior art, the broadcast/multicast boundaries of the virtual LAN and of the VNETs are equal. Thus, as mentioned above, multicast/broadcast traffic for IPX networks will be received and processed by nodes which are members of an IP subnet, if both nodes fall in the same virtual LAN.

Prior art techniques have arisen to divide networks into several virtual LANs. United States Patent No. 5,394,402 to Ross describes a virtual LAN architecture in a network which includes a backbone using a synchronous transfer mode (ATM) switching. The virtual LAN groupings act to limit the size of the multicast/broadcast domains by constraining the layer two addressing within the virtual LAN, and thus help manage the amount of multicast/broadcast packets which must be handled by a user of the network.

To cross virtual LAN boundaries, internetworking devices providing layer three routing functions are required. Thus, when a change is made in a network having a number of virtual LANs, such as a new node being added, or a user moving from one LAN segment to another LAN segment in a different virtual LAN, the VNETs must be reconfigured for the new or moved node, such as by assigning a new layer three address to the node and the like. This complication has effects throughout the network, as the internetworking devices in the system need to learn the new information, and to learn that the old information in the case of a moved node, is obsolete. Further, individual users of the virtual LANs which may have cached the old layer two MAC address of the moved node, will lose track of the node, as it will not be able to send a packet across the virtual LAN boundary with the cached layer two MAC address. Also, the use of several virtual LANs within an organization, may place constraints on layer three network definition. For instance, the IPX

network number used in the VNET of a first virtual LAN should not be used in the VNET of a second virtual LAN, because if a node moves from the first to the second, the moved node might erroneously access resources in the VNET of new virtual LAN with the network number of the VNET in old virtual LAN.

Thus, dividing a network into a number of virtual LANs, while having some benefits, also introduces complexity which may offset the benefits. It is desirable therefore to provide a more elegant method for managing traffic in large virtual LANs.

SUMMARY OF THE INVENTION

According to the present invention, network traffic management is achieved based on automatically setting up a plurality of VNETs within a single large virtual LAN. Multicast/broadcast traffic is confined to the VNET of the source, without imposing constraints on layer two addressing within the virtual LAN. Thus, when a node is moved within the network from one segment to another, it remains within the same virtual LAN, so that it may keep its layer three address or addresses, and unicast packets addressed to it from other users of the virtual LAN find their destination. Furthermore layer three network configuration is unconstrained.

The present invention can be characterized as a method for managing traffic in a network based on a set of local area network segments interconnected as a virtual LAN, and in which nodes on respective LAN segments in the set are members of VNETs. The method includes tunneling a multi-destination packet originating from a particular node in the virtual LAN, encapsulated, or otherwise reformatted, as a single destination message to a virtual net server. The virtual net server translates the multi-destination packet to a plurality of directed messages identifying nodes authorized to receive multi-destination packets from members of the particular VNET which originated the packet. The directed messages are then forwarded from the virtual net server to the authorized nodes. This way, multi-destination packets, such as

advertisement or discovery packets, are confined to a single VNET. By confining the multi-destination packets to a single VNET, unicast packets generated within the virtual LAN are then also naturally confined to that VNET. Packets are naturally confined to the VNET, because the advertisement of their address, and the procedures used to discover the addresses of others, are prevented from exiting the VNET of the particular node which issues the multi-destination packet. The present invention elegantly controls proliferation of multicast/broadcast traffic in large virtual LANs and confines unicast traffic to the VNET of the source, without introducing the complexities of prior art techniques to divide large virtual LANs into several smaller ones.

According to this aspect of the invention, the virtual net server automatically configures itself in response to the multi-destination packets received at the virtual net server, and in response to the layer three networks set up in the virtual LAN. Thus, when a virtual net server receives a multi-destination packet, it determines a virtual net domain based on the layer three network protocol and logical network which originated the packet, and the source medium access control (MAC) address of the packet. If a packet is received from a source node which had not previously sent a packet using the identified logical network identifier, then a connection is set up between the virtual net server and that source node, adding the new node to the appropriate virtual net domain. Thus, the virtual net domain is defined as a group of nodes intended to receive multi-destination packets from members of a particular VNET determined by a layer three network protocol/network identifier.

The present invention is particularly suited to connectionless backbone networks, such as FDDI, Ethernet or Token Ring LANs, using either a centralized or distributed virtual net server. In the centralized embodiment, the virtual net server is contained in the backbone LAN. In a distributed embodiment, edge devices provide management of the multi-destination packets.

According to one aspect of the invention, it can be characterized as a method for managing traffic in the network which includes a set of local area network segments, a connectionless backbone, and a plurality of edge devices which interconnect the set of LAN segments with the backbone. The method according to this aspect, and using a centralized server, includes:

detecting in an edge device on an originating LAN segment, a multi-destination packet in response to a medium access control MAC address in the multi-destination packet;

supplying the multi-destination packet from the edge device as a single destination message in the backbone to a virtual net server;

determining in response to the multi-destination packet, in the virtual net server, the virtual network of the source of the packet, and producing a plurality of single destination messages identifying nodes authorized to receive multi-destination packets from members of the determined virtual network;

forwarding across the backbone the plurality of single destination messages to edge devices coupled to LAN segments through which the authorized nodes are accessible; and

supplying the multi-destination packet from the edge devices receiving the single destination messages from the virtual net server, to LAN segments, other than the originating LAN segment, through which authorized nodes identified in the respective single destination messages are accessible.

The plurality of single destination messages are composed by encapsulating the multicast packet in single destination packets addressed to agents in each edge device through which nodes that are members of the particular virtual net domain are accessible. In the edge device, the virtual channel connection is mapped to ports of the edge device through which nodes are accessible that are members of the particular virtual net domain associated with the virtual channel. This mapping may be done in response to the source address of the multi-destination packet during the configuration process.

Alternative systems may use other types of backbone networks, including connection-oriented backbone networks.

Accordingly, a low-cost virtual LAN/virtual net (VLAN/VNET) architecture has been provided. Edge devices operate at layer 2, based on MAC address filtering. The layer 3 multi-protocol complexities are confined in the virtual net server on the backbone LAN. However, the layer 3 multi-protocol complexity only includes components necessary to decode and forward the multi-destination frames. Furthermore, the virtual NET server, the edge devices and adapters automatically learn virtual net domains of LAN segments and nodes in the system.

Unicast frames are relayed at layer 2 and automatically stay within their appropriate virtual net domains, by the inherent control of address advertisement and discovery procedures and the like.

The present invention greatly improves flexibility of network architectures by managing the flow of traffic within virtual LANs. The invention allows the creation of a plurality of VNETs within the virtual LAN according to guidelines unique to each installation, such as shared access to services, using existing logical network constructs of standard layer three protocols.

Other aspects and advantages of the present invention can be seen upon review of the figures, the detailed description and the claims which follow.

BRIEF DESCRIPTION OF THE FIGURES

Fig. 1 provides a conceptual overview of a network configured with virtual LAN domains and virtual net domains according to the present invention.

Fig. 2 is a schematic diagram of a network implementing the VLAN/VNET architecture of the present invention across an ATM LAN emulation backbone.

Fig. 3 is a schematic diagram of a network implementing the VLAN/VNET architecture of the present invention across a connectionless LAN backbone.

Fig. 4 is a functional diagram of the virtual net server and virtual net agent system used according to the architecture of Fig. 3.

Fig. 5 is a flow chart of the basic tunnelling process used in the system of Figs. 3 and 4.

Figs. 6 and 7 illustrate encapsulation of multicast packets for the tunnelling process.

Fig. 8 is a flow chart illustrating the process executed by the virtual net server and agent of Fig. 3 with a centralized server.

Fig. 9 is a flow chart illustrating the configuration routine referred in the flow chart of Fig. 8.

Fig. 10 is a flow chart illustrating a process executed by a distributed virtual net server, located on an edge device for a packet received from a LAN segment.

Fig. 11 is a flow chart illustrating the process executed by the distributed virtual net server of Fig. 10, when executed in response to a packet received from the backbone side of the edge device.

DETAILED DESCRIPTION

Fig. 1 provides a conceptual overview of a network in which the present invention operates. The network includes a plurality of LAN segments coupled to end systems or nodes on the network. The LAN segments include segments 10 - 17 which are connected to an edge device 18, and segments 19 - 26 which are coupled to edge device 27. A backbone network 28 is coupled to each of the edge devices 18 and 27 to provide interconnection among the LAN segments. Also coupled to the backbone network, may be adapters, such as adapters 30 and 31 which connect directly to end systems. Within the wired network which includes the LAN segments, edge devices, adapters and the backbone, a virtual

LAN domain 35 may be established as a layer two construct. Many virtual LANs may be implemented using higher layer procedures, such as described in the Ross patent referred to above or otherwise, but the example of one virtual LAN is used to illustrate the present invention.

5 According to the present invention, multiple virtual net domains, including virtual net domain C, virtual net domain B and virtual net domain A are set up within a single virtual LAN domain 35. A virtual net domain is defined as the set of LAN segments/ATM systems that are members of the same network layer protocol logical networks which are identified by a unique
10 network layer identifier, and may be extended to include other nodes intended to receive packets from members of this logical network.

 Virtual LAN domains contain numerous interconnected LAN segments, each with one or more attached systems (desktops, servers, routers, etc.) interconnected across a backbone 28. The utilized protocol stacks within the
15 network (e.g. IP, IPX) must be able to function properly within the virtual LAN domain.

 A virtual network configuration is utilized when variant network layer protocols, and logical networks are used within the virtual LAN. For example, a single virtual LAN wide virtual net domain may be created for IP, while
20 requiring creation of several IPX virtual net domains. Each LAN segment end system can be then individually attached to differing IPX VNETs based on policies such as desired services. Also, separate VNET domains may be created for many IP subnets and many IPX networks. Each LAN segment end system can be then individually attached to both IP and IPX VNET domains based on
25 policies such as desired services. Within a single virtual LAN domain, LAN segments and end systems may attach to multiple VNET domains.

 According to the present invention, the flow of multicast /broadcast MAC frames are kept within the associated VNET domain. All unicast MAC frames are sent across the standard virtual LAN. However, the interesting point is that
30 each unicast frame will be addressed at layer 2 to stay within its VNET domain

automatically. This means the edge devices are fast, inexpensive and simple while the VLAN/VNET architecture is optimized.

Fig. 2 illustrates the architecture of a network using ATM LAN emulation backbone with the VLAN/VNET architecture of patent application of which this is a continuation-in-part. The ATM Forum specifies a so-called LAN emulation mode by which LAN segments, and end systems directly attached by ATM adapters to the backbone, are interconnected across an ATM backbone establishing a so-called emulated LAN in the ATM environment. LAN Emulation Over ATM Specification, published by the ATM Forum, LAN Emulation SWG Drafting Group.

According to the ATM LAN emulation specifications, multi-destination packets are detected at edge devices and forwarded to a broadcast and unknown server (BUS) on the ATM backbone. The BUS takes the multi-destination packet received from an edge device, and generates a number of virtual channels to forward the multi-destination packet from the BUS, one virtual channel to each of the nodes in the emulated LAN. According to this technique, it is insured that, from the point of view of the sender of the packet, the multi-destination packet is fanned out across the point-to-point ATM backbone, like a broadcast or multicast frame would have been transmitted across a standard connectionless LAN backbone.

As can be seen in Fig. 2, an ATM LAN emulation backbone 100 is provided which is coupled to a first edge device 101 and a second edge device 102. Edge device 101 includes N ports P1-PN as illustrated in the figure. Each of the ports is coupled to a corresponding LAN segment executing a connectionless protocol, such as specified IEEE 802.x standards or other protocols like the ANSI standard Fiber Distributed Data Network (FDDI). For instance, a carrier sense multiple access with collision detect protocol such as specified in 802.3 (also called Ethernet), and a token ring protocol such as specified in 802.5 may be coupled to various ports of the edge device 101. In the same manner, edge device 102 has plurality of ports P1-PN as shown in the

figure coupled to a variety of LAN segments. With the ATM LAN emulation backbone, an ATM adapter 103 and an ATM adapter 104 may be coupled directly to the backbone. The adapters 103 and 104 are coupled directly to ATM end systems.

5 As represented by the cloud 105, a variety of other edge devices and ATM adapters may be coupled to the LAN emulation backbone 101, to establish a virtual LAN over a wide variety of LAN segments and across wide area links.

According to the present invention, a virtual net server 106, such as in an improved BUS (Broadcast and Unknown Server) in a directly attached ATM
10 end system or in an ATM switch, is coupled to the backbone 101. Also, virtual net agents 107 and 108 are implemented in the edge devices 101 and 102.

When a multicast frame is detected on a LAN segment in an edge device 101, the multicast packet is forwarded to the virtual net server 106 across the LAN emulation backbone in the manner that such packets are forwarded to the BUS.
15 Virtual net server 106 translates the multi-destination packet into plurality of directed messages which are sent across virtual channels to the virtual net agents in the edge devices, such as agents 107 and 108. The virtual net agents 107, 108 then forward the multi-destination packet out ports of the edge device on which nodes authorized receive the multi-destination packet are found. When there is
20 one user on each LAN segment, the multi-destination packet can be delivered exclusively to members of the virtual net domain.

Fig. 3 illustrates the architecture of a network using virtual net architecture according to the present invention with a backbone network which may be implemented using a connectionless protocol such as FDDI, Ethernet or
25 Token Ring. Thus, as can be seen in Fig. 3, a backbone network 120 is coupled to a first edge device 121 and a second edge device 122. The edge device 121 includes ports P1-PN as illustrated in the Fig. Each of the ports is coupled to a corresponding LAN segment executing a connectionless protocol such as specified in 802.X standards or other protocols. Alternatively, one or more

ports may be connected to an ATM edge device which extends the virtual LAN across an ATM emulation backbone.

In the same manner, edge device 122 has a plurality of ports P1-PN as shown in Fig. 3, coupled to a variety of LAN segments. Also, an end system
5 123 may be coupled directly to the backbone network 120. As represented by the cloud 124, a variety of other edge devices and end systems may be coupled to the backbone 120, to establish a virtual LAN over a wide variety of LAN segments and across wide area links.

According to the present invention, a virtual LAN and virtual net server
10 125 (VLAN/VNET server) is coupled to the backbone network 120, such as in an end system on the backbone network, or in a network intermediate system device like a router, bridge or switch on the backbone network 120. Also, virtual net agents 127 and 128 are implemented in the edge devices 121 and 122 respectively. When a multicast frame is detected on a LAN segment in an edge
15 device 121, the multi-cast packet is tunneled to the VLAN/VNET server 125 through the backbone network. The server 125 translates the multi-destination packet into a plurality of tunneled messages which are sent to virtual net agents 127 and 128 in the edge devices coupled to the backbone network 120. The virtual net agents 127 and 128 then forward the multi-destination packet out
20 port of the edge device on which nodes authorized to receive the multi-destination packet are found. When there is one user on each LAN segment, the multi-destination packet can be delivered exclusively to members of the virtual net domain using this architecture.

Fig. 4 provides a functional diagram of the virtual net server and virtual
25 net agent used according to the system of Fig. 3. Thus, an edge device 200 is illustrated in the figure. A server 201 is coupled to the edge device 200 across the LAN backbone.

The server 201 includes a decoder 203, and a plurality of virtual net tunnel
modules 204, 205, 206, 207. The edge device 200 includes an agent 208 which
30 operates with the server 201.

In operation, an edge device 200 receives on an incoming port 210 from an originating LAN segment a multicast or broadcast packet. This packet is then routed using tunneling 211 under control of the agent 208 to the server 201.

5 The decoder 203 in the virtual net/virtual LAN server 201 (V/V server) determines the virtual net domain of the packet in response to the network protocol of the multicast /broadcast packet and the network identifier, if used in the identified protocol, by layer three protocol constructs in the packet. It then passes the multi-destination packet to the appropriate virtual net tunnel module. A virtual net domain exists for each network address value (for example each IP
10 subnet value) supported by a given network protocol. When the frame does not contain a network identifier (for instance a NetBIOS frame) only one virtual net tunnel module exists for it in the server 201. Thus, if the multicast /broadcast packet is an IPX packet, then it is forwarded across either line 212 or line 213 to the VNET tunnel module 206 or module 207 for corresponding network
15 identifiers. If the multicast /broadcast packet is an IP protocol packet, then it is forwarded across line 214 or line 215 to one of the VNET tunnel modules 204 or 205 for corresponding subnets. For the purposes of this example, the packet is passed to the subnet 1 tunnel module 204. The subnet 1 tunnel module 204 includes a table 216 which maps the virtual net domain to established tunnels
20 in the backbone. Established tunnels provide direct paths (e.g. 217) by means of single destination packets to agents in edge devices 200 on which ports authorized to receive the frame are found. Thus, in this example, the agent 208 in the edge device 200 receives the multicast packet across tunnel 217 and forwards the packet out the appropriate ports. The agent includes table 218
25 which maps the tunnel on which the multicast packet is received, to ports on which nodes authorized to receive the packet are found, using the source MAC address of the packet to make sure that it is not sent back on the originating LAN segment 210. In the illustrated example, the packet is sent on port 219 and port 220 by the agent 208, but not on other ports of the edge device and not on
30 the port coupled to originating LAN segment 210.

Alternatively, a VNET encapsulation can be utilized which provides the ability for each of the virtual net tunnel modules to share tunnel addresses established for each of the edge devices. However, the agent at the edge device must be able to handle the decapsulation of the frame as it received and route it appropriately.

Standard end system adapters can utilize this virtual net architecture. The configuration steps are not required for end systems because there is no edge device connected to them which operates as a proxy for other LAN segments. Since the directly connected end systems do not register as proxies within the server, the end system is automatically distinguished from an edge device. Thus, only one tunnel per virtual LAN is set up from the server to the end system. All associated virtual net membership entries set the tunnel identifier values to the same tunnel for the adapter. In the configuration process, the virtual net membership entry, layer two MAC address is not forwarded to the end system adapter, because it is not necessary for use there.

Unicast frames are forwarded using the standard LAN processing mode, and are naturally confined their own virtual net domain. Optionally, a unicast privacy checking algorithm can be added, by having the edge devices check with the V/V server 201, the first time they see a destination value from a LAN segment. From the membership lists in the server 201, it can be verified that the source and destination addresses remain within the same virtual net domain.

The basic tunneling process according to the present invention is illustrated with respect to Fig. 5, 6 and 7. In Fig. 5, the basic process is described, which begins with receiving an incoming multicast packet at a virtual net agent on an edge device. The edge device forwards the multicast packet on attached segments, with or without filtering by virtual net domain (block 150). Alternatively, it may defer forwarding the packet on attached segments until it receives the multicast back from the server.

Next, the virtual net agent encapsulates the multicast packet and tunnels it to the virtual net/virtual LAN server on the backbone network (block 151). At

the virtual LAN/virtual net server, the multicast packet is decapsulated from the tunneled message, and encapsulated in a new tunnel packet for forwarding to other virtual net agents (block 151). At the agents, the tunneled messages are decapsulated, and the outgoing multicast packets is forwarded on attached
5 segments, other than the segment which originated the message and segments which already received the message (block 153).

The tunneling process can be understood with references to Fig. 6 and 7. Basically, tunneling involves encapsulating a multi-destination packet in a single destination packet having a MAC address of the destination of the tunnel,
10 and a source address equal to the source of the tunnel. Thus, the agent in the edge device will encapsulate the message as illustrated in Fig. 6 where the multicast frame 155 is encapsulated in a single destination packet having the server address 156 as a destination address and other supporting control fields, such as the frame check sequence 157, surrounding the multicast frame 155. At
15 the receiving end of the tunnel, the server receives the frame and processes it. It discovers that the frame is a tunneled multicast frame, and using the process described with respect to Fig. 4, encapsulates the frame in a tunnel directed to the agent as shown in Fig. 7. Thus, a tunnel from the server to the agent will carry destination address equal to the agent address 160, the multicast frame 155
20 will be encapsulated within the packet. Supporting control fields, such as the frame check sequence 161 and the like, are included within the tunneled packet.

The tunneling process can take a variety of formats. For instance, each tunnel may be established by setting up a specific destination and source address for each tunnel handler/edge device pair. This way, the agent and server must
25 maintain a number source and destination addresses, and correlate those with specific VNETs. Alternatively, a single address may be used for the server and a single address used for each agent, and the tunnel packet will carry control fields which specify the information needed to recognize the packet as a tunneled packet.

Fig. 8 provides a flow chart illustrating the handling of multicast and broadcast packets according to a centralized server embodiment of this system. The process begins when an edge device receives a frame (block 250). The device determines the type of frame (block 251). If frame is a unicast frame, then it is handled with standard LAN techniques (block 252). If the frame is a multicast or broadcast, then the edge device forwards the frame using tunneling to the V/V server (block 254).

In the server, the decoder determines the virtual net domain of the frame (block 255). Then the frame is passed to the determined virtual net tunnel handler (block 256). The virtual net tunnel handler determines whether the source of the multicast packet has a corresponding entry in its membership list (block 257). If an entry is found, then the frame is forwarded as a tunneled message as set up in the membership list (block 258).

If no entry was found in the membership list at the virtual net tunnel handler in the test of block 257, then an automatic configuration routine is executed (block 260). After the configuration routine, then the process proceeds to block 258 to forward the frame across the established tunnels for the virtual net tunnel handler.

The process of block 258 results in edge devices receiving the multicast/broadcast frame. Each edge device which receives the multicast/broadcast frame, then sends the frame once on ports to members of the virtual net domain. This is done by the edge device maintaining a table which maps the tunnel on which the frame is received to specific ports, or modules accessible through the ports, of the edge device. However, the edge device does not send the multicast packet back out on the segment which originated the packet. This is determined by checking the source address of the multicast/broadcast frame, and comparing that source address with the address of devices on the respective ports (block 259).

Fig. 9 illustrates the configuration routine executed at block 260 of Fig. 8. According to this routine, the virtual net tunnel handler sets up a tunnel from the

virtual net tunnel handler to the originating edge device (block 270). An entry is created in the virtual net membership list for the source of the packet (block 271). The entry includes a source MAC address of the originating end station and a tunnel identifier (i.e. MAC address of the agent in the edge device) (block 272). After creating the entry, the virtual net tunnel handler sends the source MAC address across the established tunnel to the originating edge device (block 273). The edge device then stores the received source MAC address in the virtual channel/virtual net membership list maintained by the agent (block 274). This MAC address is utilized to map incoming frames on this tunnel to the appropriate ports of the edge device.

Utilizing the process of Figs. 8 and 9, each virtual net tunnel handler 204, 205, 206, 207 as shown in Fig. 4, establishes a tunnel to each edge device which includes a port through which a member of the virtual net domain is found. These established tunnels provide a mechanism for distributing the multiple destination packets efficiently across the backbone. The decoder in the server 201 maps the incoming packet to the virtual net tunnel handler which maps the frame based on a membership list to a set of established tunnels. The edge devices map frames incoming on specific tunnels to ports of the edge device. This tightly controls the propagation of multiple destination packets within the appropriate virtual net domain of the originating device.

As mentioned above, the V/V server can be distributed to the edge devices, rather than executed in a centralized site. Figs. 10 and 11 illustrate a process which is executed in the edge devices according to this distributed virtual net server model. Thus, in Fig. 10, the process executed by the edge device when it receives a packet from the LAN segment on the user side of the edge device is shown. The process begins with receiving a packet from the user side (block 300). The process then determines the type of frame (block 301). If it is a unicast frame, then it is handled with standard LAN procedures (block 302). If it is a multicast or broadcast packet, the distributed virtual net server determines the VNET domain of the frame (block 303). The source address of

the multicast packet is added to a VNET domain list which is maintained in the edge device, if it is not already there (block 304). Finally, a multicast packet is forwarded using tunneling to other edge devices, and sent on local LAN segments. Optionally, transmission on the local LAN segments can be filtered by VNET membership (block 305).

The process shown in Fig. 11 is executed by edge devices receiving a packet from a tunnel from other edge devices. Thus, the process begins with receiving a frame from the tunnel on the backbone network (block 310). The edge device next determines the type of frame (block 311). If it is a unicast frame, it is handled with the standard LAN processes (block 312). If the frame is a multicast or a broadcast packet, the edge device determines whether the packet has a local source address (block 313). If it has a local source address, then it is discarded (block 314), because the process described in Fig. 10 has already forwarded the packet to the local LAN segments. If the packet does not have a local source address, then the edge device sends the frame once on ports coupled to members of the virtual net membership list which matches the multicast packet (block 315). Accordingly, the virtual net server can be distributed to the edge devices in the virtual LAN architecture.

The present invention provides management of traffic in a virtual LAN environment according to the concept of a virtual net domain. To maintain virtual net domain boundaries, edge devices operate at layer 2, while limited layer 3 complexity is centralized in, for example, an improved server, or alternatively in a distributed virtual net server. The virtual net server, the edge devices and adapters automatically learn virtual net domain membership among nodes on connected LAN segments.

The foregoing description of a preferred embodiment of the invention has been presented for purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Obviously, many modifications and variations will be apparent to practitioners

skilled in this art. It is intended that the scope of the invention be defined by the following claims and their equivalents.

What is claimed is:

CLAIMS

1. For a network including a set of local area network (LAN) segments interconnected as a virtual LAN, in which nodes in the virtual LAN are members of one or more logical networks, a method for managing traffic in the network, comprising:

detecting a multi-destination packet on a LAN segment within the set;
determining, in response to the multi-destination packet, the logical network for which the detected multi-destination packet is intended; and
tunneling the multi-destination packet to nodes authorized to receive multi-destination packets intended for members of the determined logical network.

2. The method of claim 1, including after said step of detecting:
tunneling the detected multi-destination packet by a single destination packet to a server; and
producing a plurality of single destination messages in said server to deliver the information carried by the multi-destination packet.

3. The method of claim 1, wherein the network includes a connectionless backbone communication path, and a plurality of edge devices which interconnect the set of LAN segments and the backbone communication path, and including:

detecting in a particular edge device, a multi-destination packet on a particular LAN segment in response to a medium access control MAC layer address in the multi-destination packet;
tunneling the detected multi-destination packet from the particular edge device across the backbone communication path to a server; and
producing a plurality of single destination messages in said server to deliver the information carried by the multi-destination packet.

4. The method of claim 3, wherein the connectionless protocol of the backbone network comprises one of a carrier sense multiple access protocol and a token passing ring protocol.

5. The method of claim 2, including:
forwarding the plurality of single destination messages to edge devices;
and
supplying in the respective edge devices, the information from the multi-destination packet to LAN segments associated with the single destination message received in the edge device.

6. The method of claim 2, including:
automatically learning the logical networks of which nodes on the set of LAN segments are members in response to the detected multi-destination packets.

7. For a network including a set of local area network (LAN) segments, a backbone path, and a plurality of edge devices interconnecting the plurality of LAN segments with the backbone path, and in which nodes on respective LAN segments in the set are members of logical networks defined in layer three or higher, a method for managing traffic in the network, comprising:
detecting in response to a medium access control MAC address of a packet received at an edge device, a multi-destination packet originating from a particular LAN segment;
determining the logical network of the multi-destination packet;
translating the multi-destination packet to a plurality of tunneled messages carrying information from the multi-destination packet;
forwarding across the backbone path the plurality of tunneled messages to edge devices; and

supplying the multi-destination packet in response to the tunneled messages from the edge devices receiving the directed messages to LAN segments through which nodes authorized to receive multi-destination packets from members of the determined logical network are accessible, except for the particular LAN segment on which the multi-destination packet originated.

8. The method of claim 7, wherein the backbone path includes a connectionless network.

9. The method of claim 8, wherein the connectionless network comprises one of a carrier sense multiple access protocol and a token passing ring protocol.

10. The method of claim 7, including;
automatically learning members of logical networks in response to multi-destination packets.

11. The method of claim 7, wherein the multi-destination packet comprises an address advertisement packet.

12. The method of claim 7, wherein the multi-destination packet comprises an address discovery packet.

13. The method of claim 7, including executing the steps of determining, translating and forwarding in a server coupled to the backbone communication path.

14. For a network including a set of local area network (LAN) segments, a LAN backbone path, and a plurality of edge devices interconnecting the plurality of LAN segments with the backbone path, and in which nodes on

respective LAN segments in the set are members of logical networks defined in layer three or higher, a method for managing traffic in the network, comprising:

detecting in an edge device on an originating LAN segment, a multi-destination packet intended for members of a particular logical network, in response to a medium access control MAC address in the multi-destination packet;

producing in response to the detected multi-destination packet, a plurality of directed messages for nodes authorized to receive multi-destination packets of particular logical network;

tunneling in the backbone path the plurality of directed messages to edge devices coupled to LAN segments through which the authorized nodes are accessible; and

supplying the multi-destination packet in response to the directed messages, from the edge devices receiving the directed messages, to LAN segments, other than the originating LAN segment, through which authorized nodes are accessible.

15. The method of claim 14, including:

tunneling the multi-destination packet from the edge device across the backbone path to a server; and

executing the steps of producing and forwarding in the server.

16. The method of claim 14, wherein the step of producing includes:

identifying the logical network of the multi-destination packet; and

composing the plurality of directed messages so that they identify nodes intended to receive multi-destination packets of the identified logical network.

17. The method of claim 16, wherein the step of composing includes:
configuring tunnel connections for each logical network to edge devices
through which nodes intended to receive multi-destination packets of the logical
network are accessible; and

5 associating the plurality of directed messages with configured tunnel
connections for the identified logical network.

18. The method of claim 17, including:
automatically configuring the configured tunnel connections in response
10 to detected multi-destination packets.

19. The method of claim 17, wherein the step of supplying the multi-
destination packets from the edge devices includes:

mapping the configured tunnel connections to ports on the edge devices
15 through which nodes authorized to receive multi-destination packet of the
identified logical network are accessible; and

transmitting the multi-destination packet, in response to the mapping and
the particular configured tunnel connection across which the directed message is
received, to ports of the edge device other than a port coupled to the originating
20 LAN segment.

20. The method of claim 19, including:
automatically configuring the configured tunnel connections in response
to multi-destination packets received at the server, including

25 detecting in the server the logical network with which the multi-
destination packet is associated and the source MAC address of the multi-
destination packet,

establishing a tunnel connection with the edge device which
detected the multi-destination packet, and

mapping in the edge device the tunnel connection to a port on which the source MAC address of the multi-destination packet is accessible.

5 21. The method of claim 14, wherein the multi-destination packet comprises an address advertisement packet.

 22. The method of claim 14, wherein the multi-destination packet comprises an address discovery packet.

10 23. The method of claim 14, including:
 learning in the edge devices the logical networks in use by nodes on respective ports of the edge devices; and wherein the step of supplying is responsive to the learned logical networks.

15 24. For a network including a set of local area network (LAN) segments, an LAN backbone path, and a plurality of edge devices interconnecting the plurality of LAN segments with the backbone path, and in which nodes on respective LAN segments in the set are members of logical networks defined in layer three or higher, a method for managing traffic in the
20 network, comprising:

 detecting in a first edge device on an originating LAN segment, a multi-destination packet intended for members of a particular logical network, in response to a medium access control MAC address in the multi-destination packet, the edge device having ports coupled to the originating LAN segment
25 and at least one additional LAN segment in the set;

 first supplying the detected multi-destination packet out ports of the first edge device, other than the port coupled to the originating LAN segment, through which nodes authorized to receive messages from members of the particular logical network are accessible;

encapsulating and forwarding from the first edge device, the multi-destination packet as a single destination packet on the backbone path to a server;

5 producing in the server in response to the detected multi-destination packet, a plurality of directed messages for edge devices on the backbone path;

encapsulating and forwarding in respective single destination packets on the backbone path the plurality of directed messages to edge devices coupled to LAN segments through which the authorized nodes are accessible; and

10 second supplying the multi-destination packet in response to the directed messages, from edge devices receiving the directed messages, to LAN segments, other than LAN segments on the first edge device, through which authorized nodes are accessible.

15 25. The method of claim 24, wherein the multi-destination packet comprises an address advertisement packet.

26. The method of claim 24, wherein the multi-destination packet comprises an address discovery packet.

20 27. The method of claim 24, including:
learning in the edge devices the logical networks in use by nodes accessible through respective ports of the edge devices; and wherein the step of first supplying and the step of second supplying are responsive to the learned logical networks.



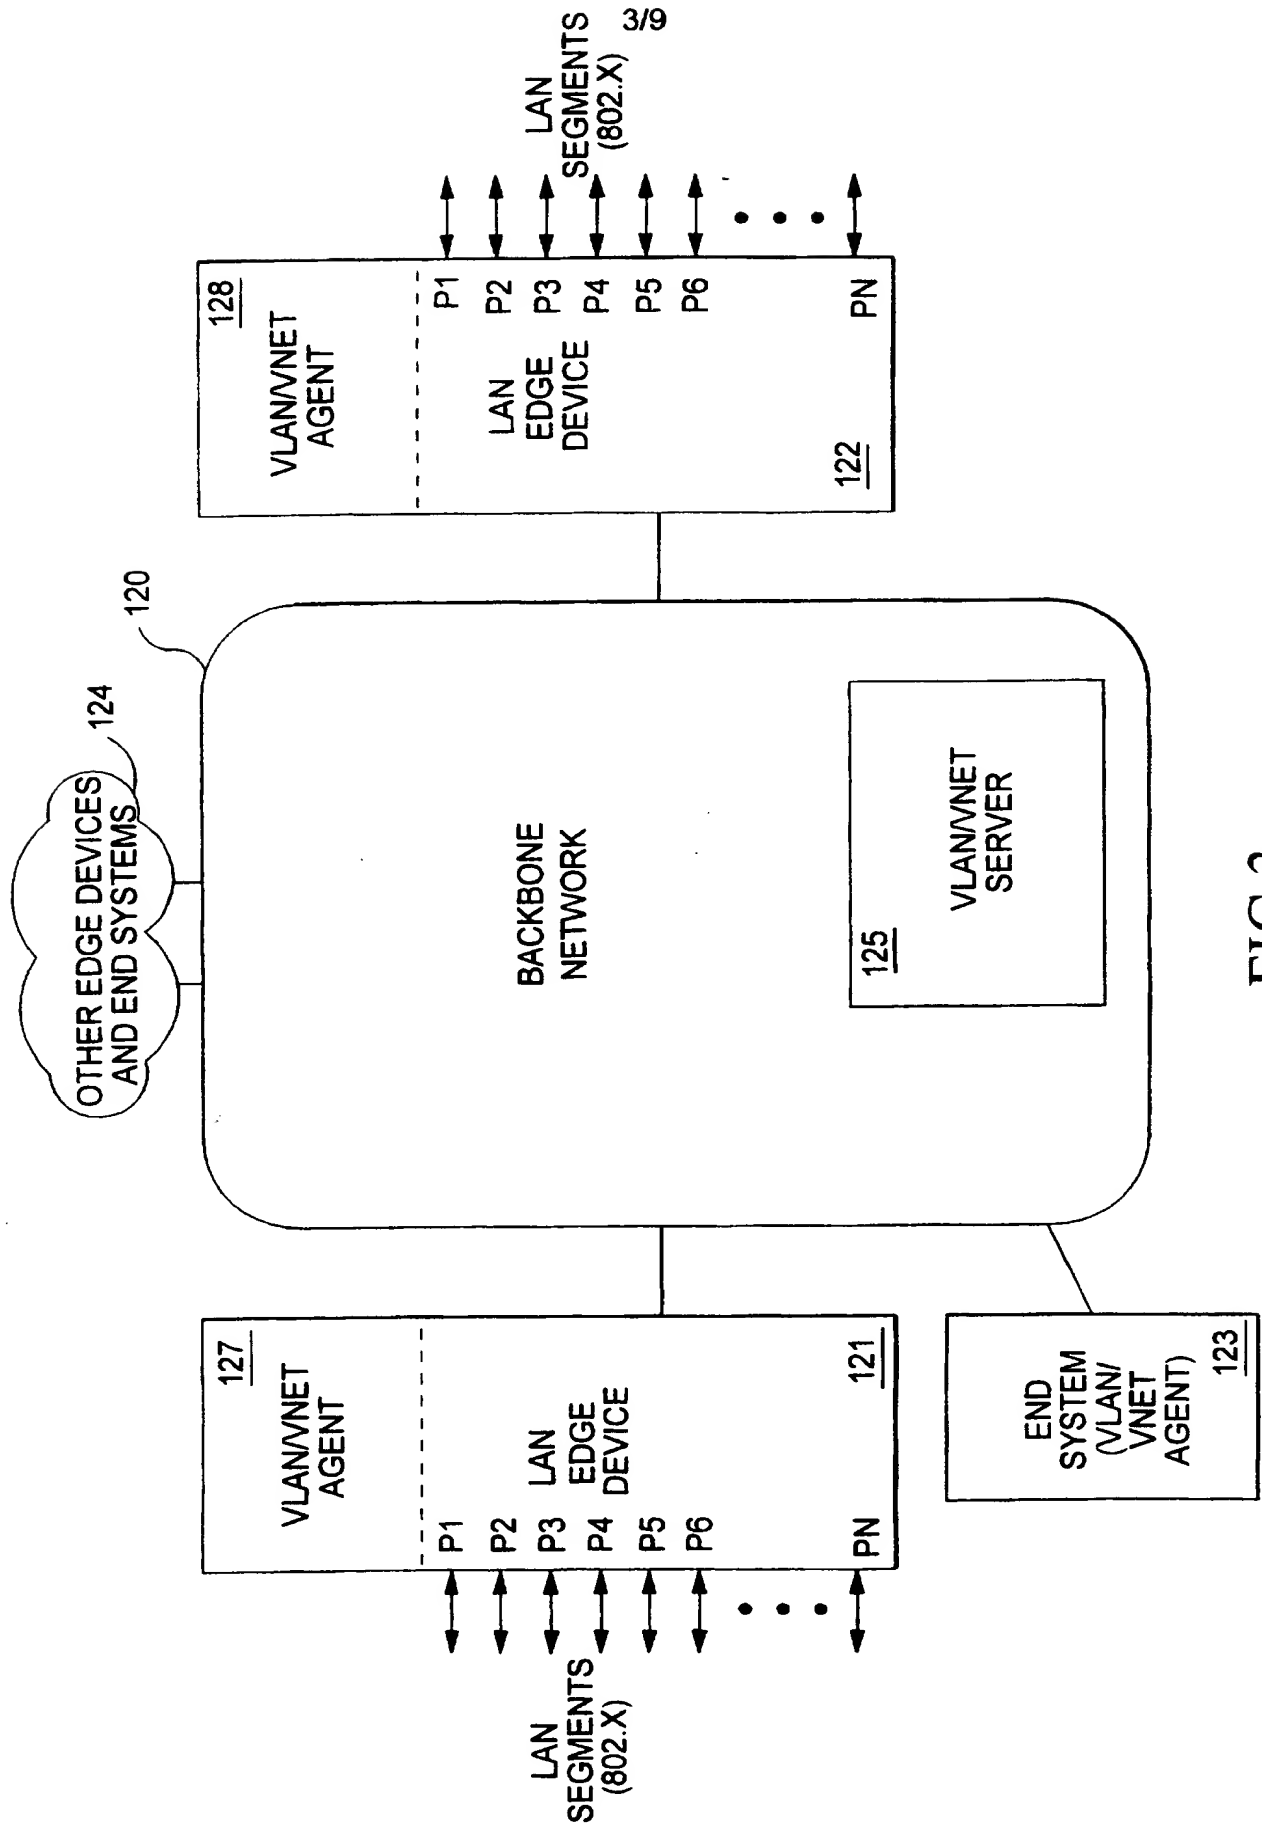


FIG.1



SUBSTITUTE SHEET (RULE 26)



SUBSTITUTE SHEET (RULE 26)




FIG.4

5/9




FIG. 5




FIG. 6




FIG. 7

6/9




FIG. 8




FIG.9

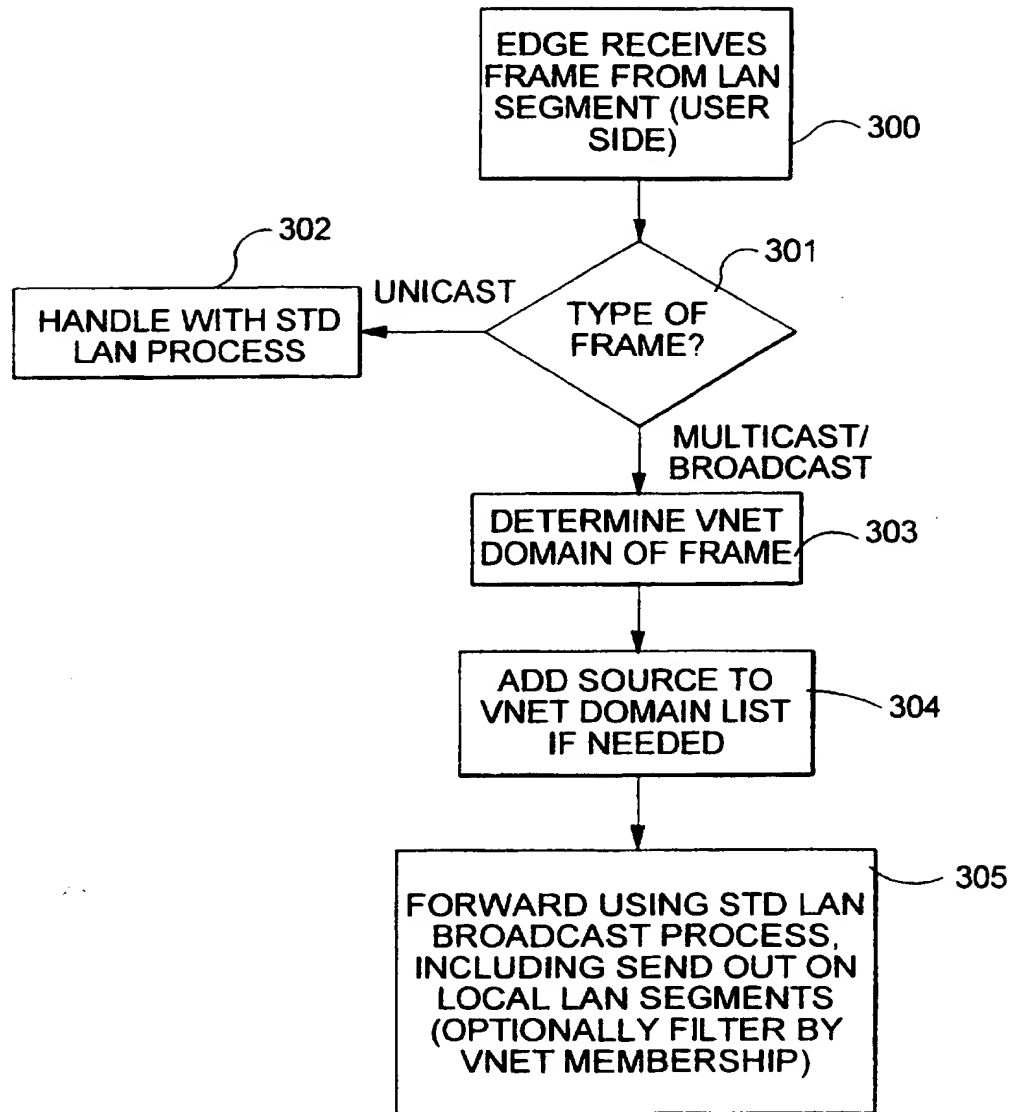


FIG.10




FIG.11

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US97/12314

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : G06F 13/00

US CL : 395/200.68, 200.79, 200.6, 200.75; 370/351, 401, 409

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 395/200.68, 200.79, 200.6, 200.75; 370/351, 401, 409

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS, RFC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 5,519,704 A (FARINACCI et al) 21 May 1996, col. 7 line 1 et seq. and column 9.	1-27
Y	SIMPSON, W. IP in IP Tunneling RFC-1853 October 1995 Whole document.	1-27
Y	SWALLOW, G., Frame Relay PVC Multicast Service and Protocol Description (FRF.7) Frame Relay Forum Implementation agreement October 1994 Whole document. Especially sections 5.1 and 9.1.2	1-27
Y	US 5,321,693 A (PERLMAN) 14 July 1994, col.3 line 17 et seq. and column 5 line 4 et seq.	1-27

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:	*T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be of particular relevance	*X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	*Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*Z* document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means	
P document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

12 SEPTEMBER 1997

Date of mailing of the international search report

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer:

MEHMET GECKIL

Telephone No. (703) 305-9676

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☒ FADED TEXT OR DRAWING
- ☒ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☒ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☒ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

THIS PAGE BLANK (USPTO)